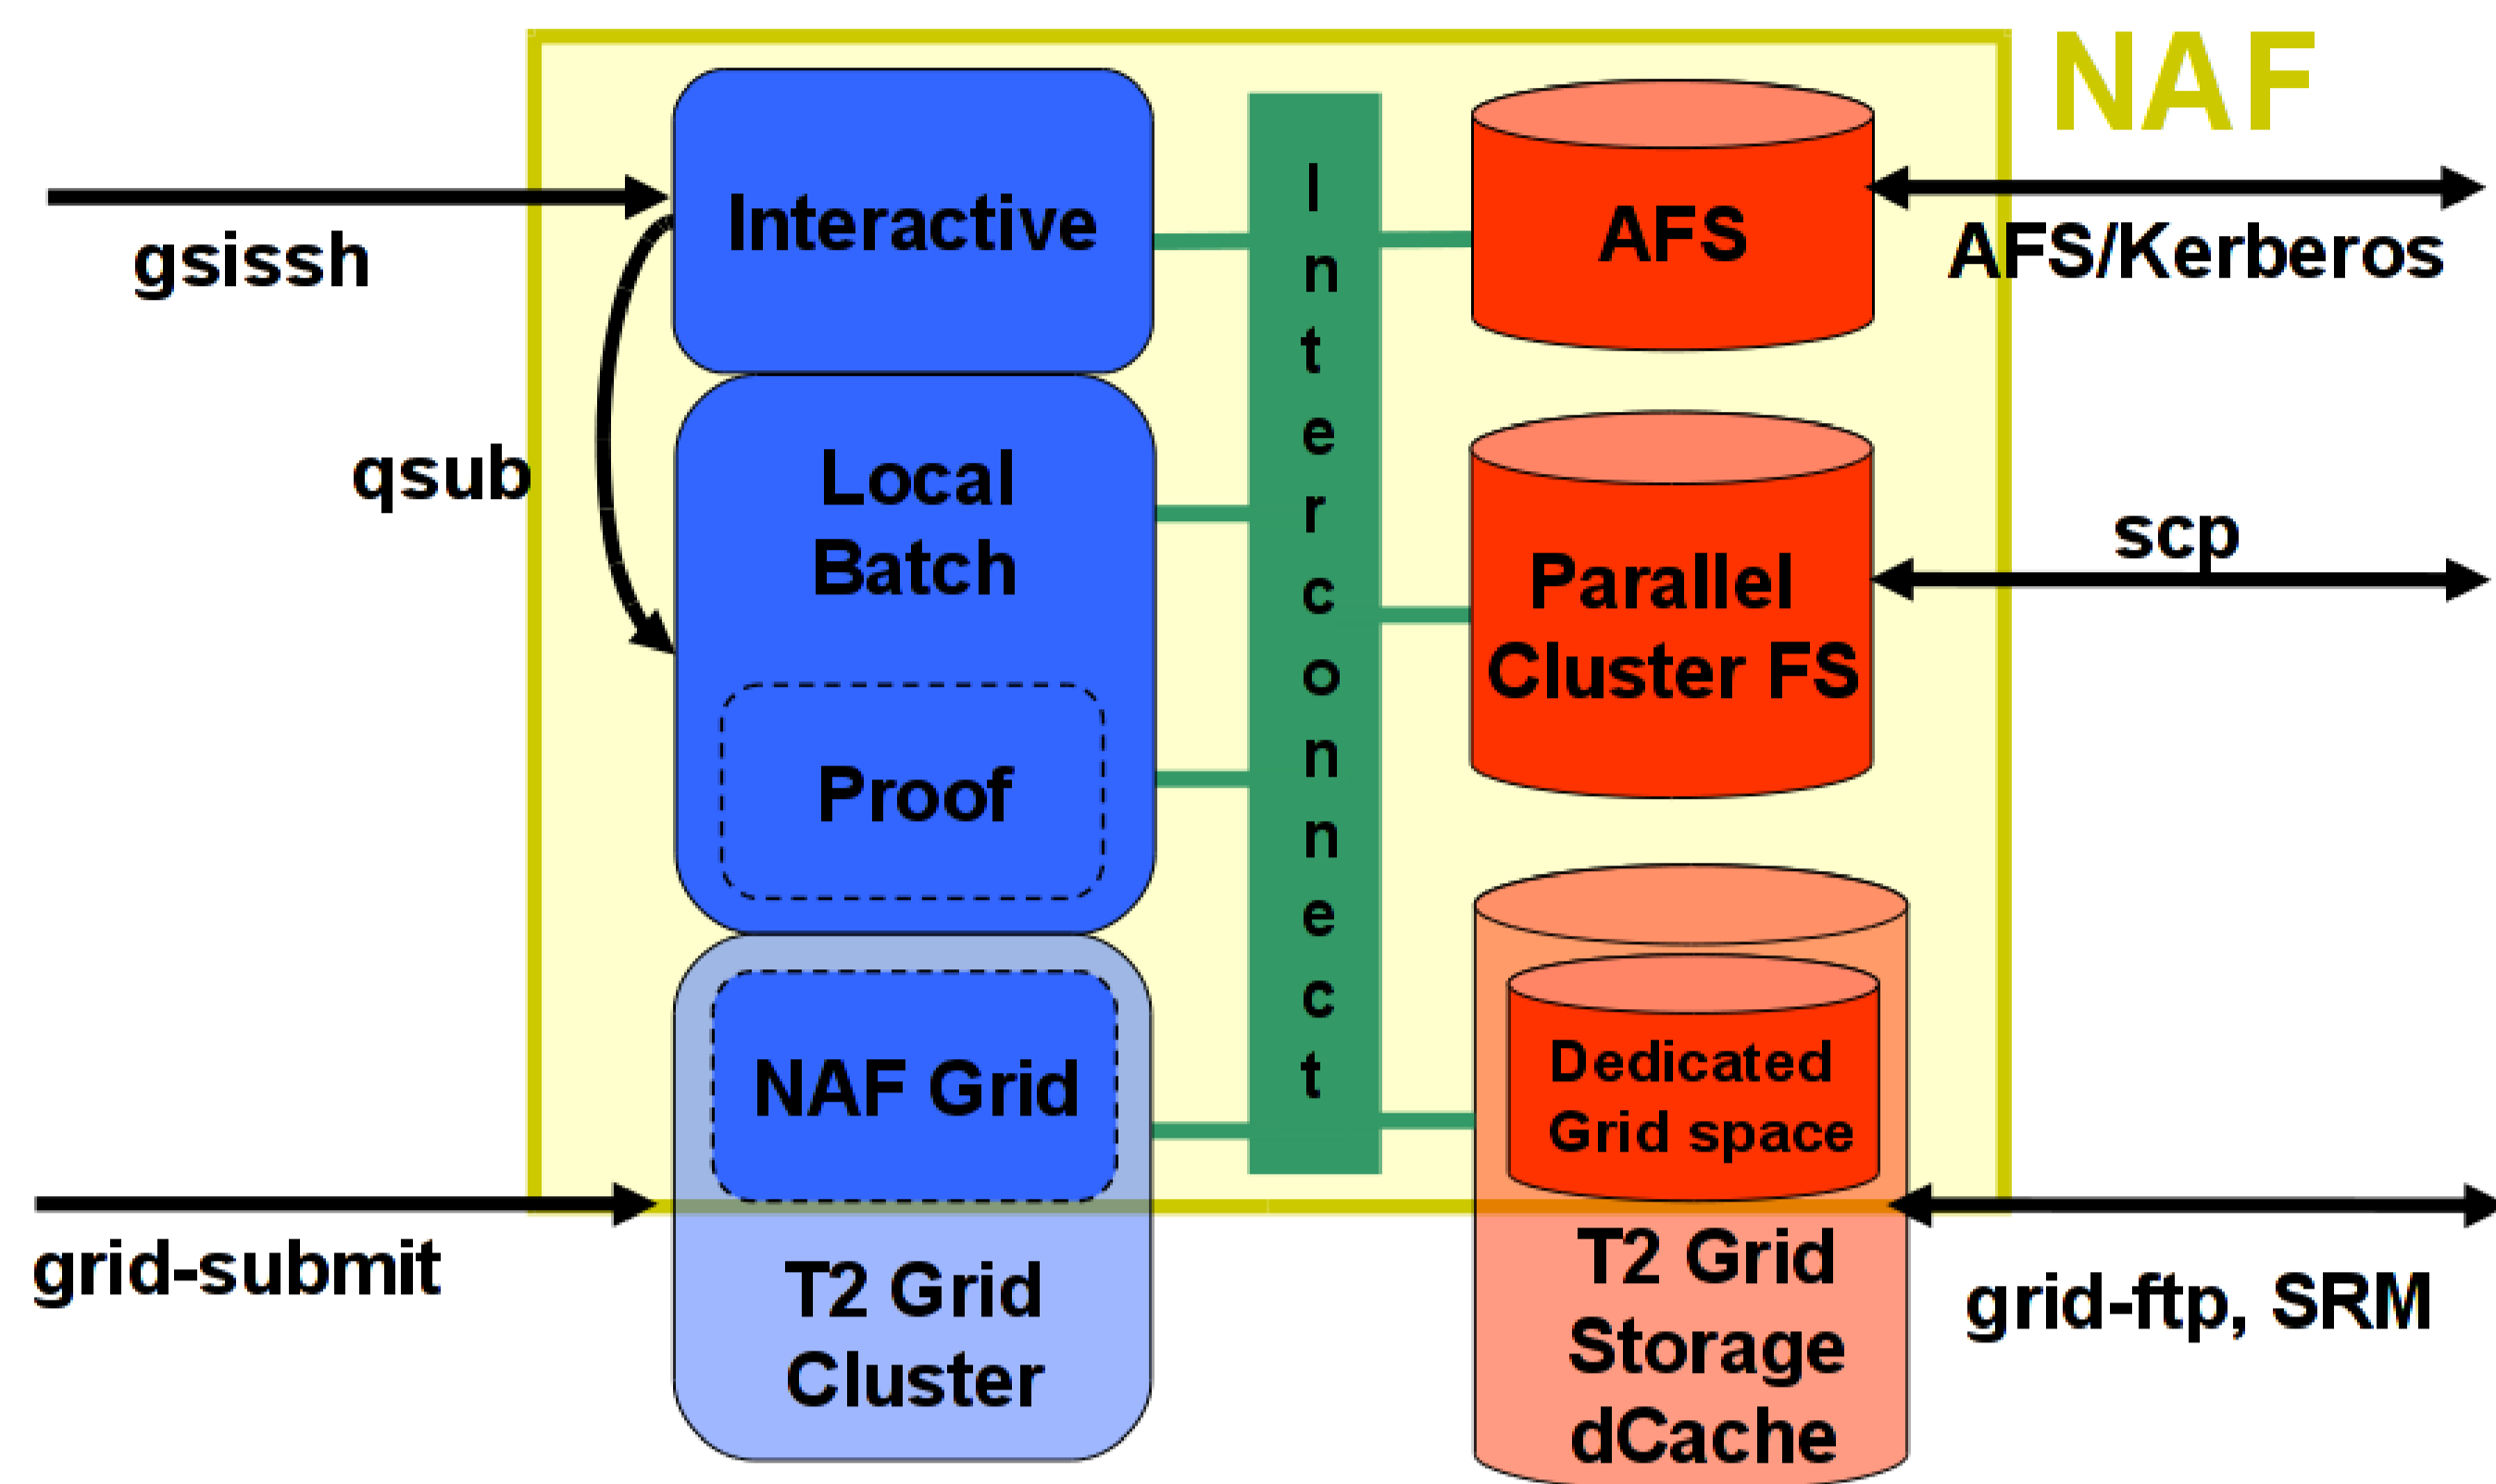


The German National Analysis Facility as a tool for ATLAS analyses.

W. Ehrenfeld, K. Leffhalm and S. Mehlhase for the ATLAS Collaboration



The National Analysis Facility (NAF)



General information

- > Set up in the framework of the Helmholtz Alliance: Physics at the Terascale
- > Analysis platform for LHC and ILC experiments
- > Exclusively for users at German institutes

General infrastructure

- > Build and operated jointly by DESY-IT/DV in Hamburg (HH) and Zeuthen (ZN)
- > Interactive session, SGE batch system and dedicated Grid element
- > AFS and Lustre storage system and dedicated Grid storage (HH & ZN)
- > For more details see NAF talk (PS05-3-330)

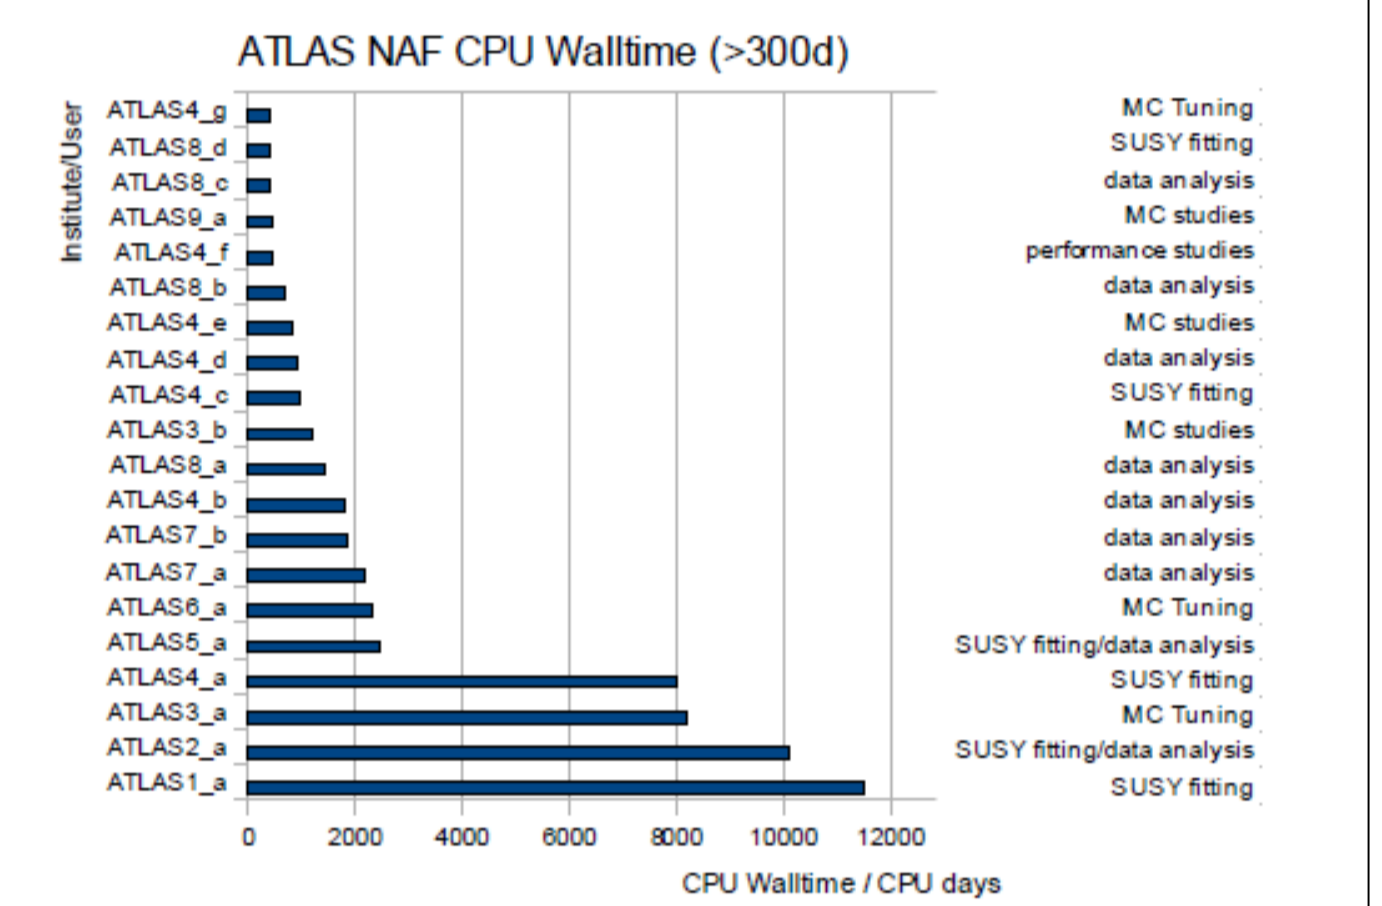
ATLAS specific infrastructure

- > 4 work group server (interactive sessions)
- > SGE batch system share of 25% out of 1400 cores including ATLAS high priority queue

Usage statistics

User statistics

- > 500 total users (300 ATLAS users)
- > 100 active ATLAS users (within last 30 days)
- > 1170 TB available data storage
 - > 740 TB T2 Grid storage
 - > 340 TB NAF Grid storage
 - > 90 TB local storage
- > AFS user space
 - > 1-5 GB AFS home directory (\varnothing 195 MB)
 - > 1-10 GB AFS scratch per user (\varnothing 620 MB)



Data availability

Data distribution policy

ATLAS physics data as well as important Monte Carlo samples are provided centrally for all users and further storage is provided for individual and group data.

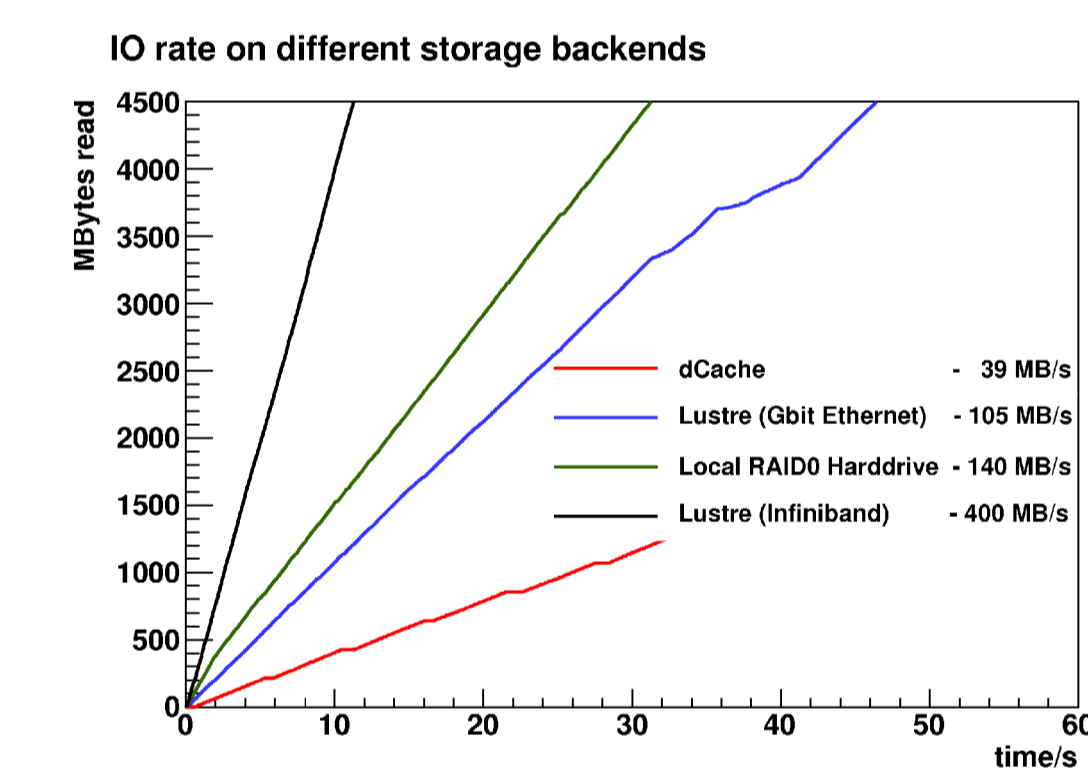
- > Full AOD dataset available at NAF: official ATLAS space tokens extended by NAF storage to extend official ATLAS data share for DESY Tier2
- > LocalGroupDisks space tokens as local Grid storage for analyses data (ntuples) and archival (older AOD versions)
- > All Grid storage included into the ATLAS distributed data management
- > Data is directly accessible via dCache from work group servers, SGE batch system and Grid
- > Fast local storage is provided using the Lustre filesystem, which offers short access times for large files and short analysis turn-around

ATLAS analysis support

Common tools

To allow for an efficient work flow common tools are provided and can be requested from DESY-IT for all users.

- > Mainly provided centrally by DESY-IT
 - > ROOT (5.26 and older)
 - > Python (2.5/2.6)
 - > Variety of text-based and X-editors
 - > LaTeX
 - > AutoProxy



Deployment of GRID tools

To exploit the full capabilities of the Grid also at the NAF, Grid-submission tools are provided centrally.

- > All recent Ganga versions are made available locally by DESY-IT in a semi-automatic fashion
- > The recent Panda client is automatically installed
- > The recent version of the DQ2 tools (distributed data management) is made available locally by DESY-IT in a semi-automatic fashion
- > Recent versions of the Grid User Interface are made available by DESY-IT
- > Setup of all tools via a simple setup program

User support

Users are supported in all ATLAS software related questions. If not NAF specific, users might be forwarded to the appropriate ATLAS help forums.

Deployment of ATLAS software

All software and tools necessary for a variety of possible analysis scenarios is provided at short time scales and centrally for all ATLAS users.

- > Recent ATLAS base software releases, official as well as working group specific caches and ATLAS data base releases are installed in a semi-automatic fashion
- > Flat ATLAS condition data are taken from the DESY Tier2 sites matching the location (HH/ZN) where the job runs
- > Access to ATLAS software and conditions data works out of the box for all users.

Workload

- > 0.2 FTEs for storage maintenance (shared between 1 persons)
- > 0.1 FTEs for software deployment (shared between 4 persons)
- > 0.3 FTEs for user administration and support (shared between 2 persons)

ATLAS analysis use cases

N-tuple production on the Grid using Ganga job submission (I)

With direct access to general and dedicated grid resources, production jobs are submitted using local installations of Ganga and Panda. Skimmed/slimmed/thinned data sets (D2PDs) or flat n-tuples (D3PDs) are created using AODs or ESDs as input and are stored at local grid storage elements. These derived format allow for shorted turn-around times in final analyses and can be accessed directly via dCache on all NAF systems.

SGE batch based analysis using data from Grid storage element via dCache (II)

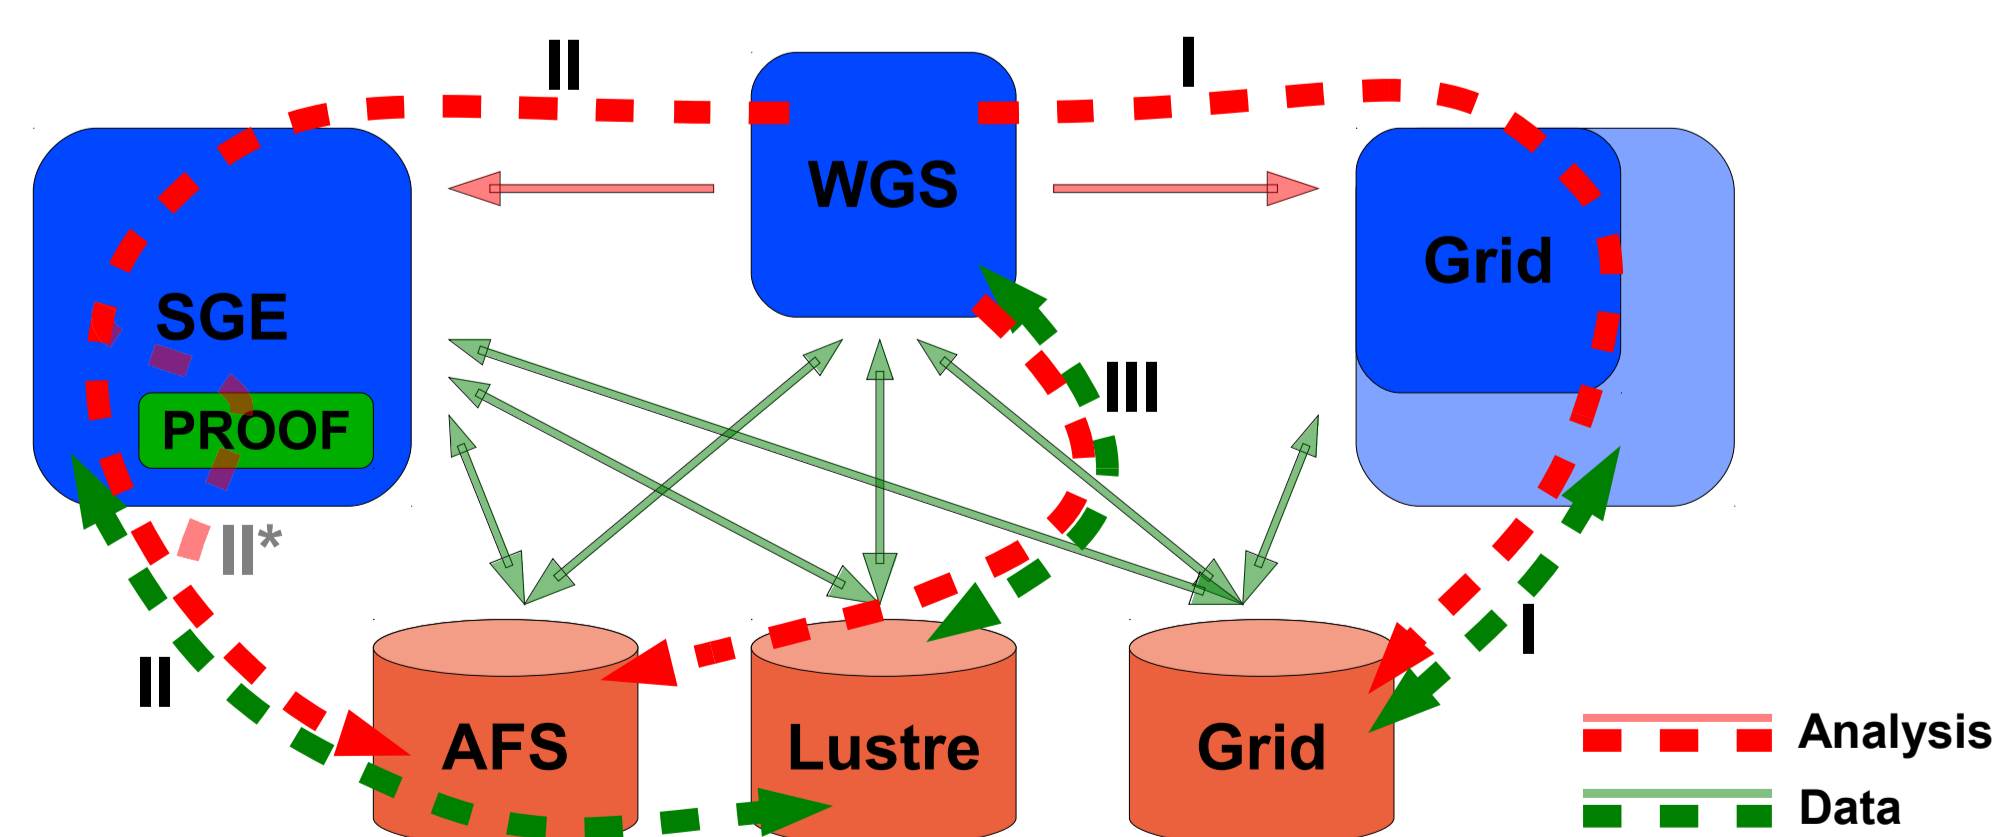
Medium sized jobs, e.g. production of n-tuples from derived formats stored locally on grid storage elements, can be performed with short turn-around times using the SGE batch system. The output can be stored e.g. on scratch space available on AFS or Lustre. Job submission and monitoring can done with standard SGE tools or Ganga.

PROOF cluster on demand on a batch system (II*)

Support of an easy to use and simple script to setup a fully working PROOF cluster on demand for a user. The script starts a PROOF master on the local work group server and submits multiple proof clients on the SGE batch system. PROOF can be used for data analysis based directly on TSelector or TSelector based frameworks as Sframe. More details on poster PO-MON-037.

Interactive analysis on a work group server using fast Lustre storage (III)

Using the fast Lustre storage element, end user analyses and the production of final plots can be carried out using interactive sessions on the dedicated ATLAS work group servers. These provide all tools necessary, such as ROOT, editors, LaTeX and more.



Operations Review

The German ATLAS NAF team has collected extensive experiences in supporting the German ATLAS users during the 3 years of NAF operations:

- > Software installation can easily be done via scripts and automated by cron jobs
- > Supporting experiment software for forward compatibility (operating system) is time consuming and does not always work out of the box
- > AFS storage management is easy
- > Local storage is always tight and users are not cleaning up on their own. An automatic cleanup based on access time is preferable
- > User Grid storage is unlimited. Efficient space administration and management is only possible within the ATLAS distributed data management system
- > User support is time consuming.